Robust and Low-Rank Representation for Fast Face Identification with Occlusions

Michael Iliadis, *Student Member, IEEE,* Haohong Wang, *Member, IEEE,* Rafael Molina, *Member, IEEE,* and Aggelos K. Katsaggelos, *Fellow, IEEE*

arXiv:1605.02266v1 [cs.CV] 8 May 2016

***Abstract*—In this paper we propose an iterative method to address the face identification problem with block occlusions. Our approach utilizes a robust representation based on two characteristics in order to model contiguous errors (e.g., block occlusion) effectively. The first fits to the errors a distribution described by a tailored loss function. The second describes the error image as having a specific structure (resulting in low-rank). We will show that this joint characterization is effective for describing errors with spatial continuity. Our approach is computationally efficient due to the utilization of the Alternating Direction Method of Multipliers (ADMM). A special case of our fast iterative algorithm leads to the robust representation method which is normally used to handle non-contiguous errors (e.g., pixel corruption). Extensive results on representative face databases document the effectiveness of our method over existing robust representation methods with respect to both identification rates and computational time.**

Code is available at Github, where you can find implementations of the F-LR-IRNNLS and F-IRNNLS (fast version of the RRC) : https://github.com/miliadis/FIRC

***Index Terms*—Face Identification, Robust Representation, Low-Rank Estimation, Iterative Reweighted Coding.**

I. Introduction

FACE IDENTIFICATION (FI) focuses on deducing a subject's identity through a provided test image and is one of the most popular problems in computer vision [1]. Typically, test images exhibit large variations, such as occlusions. Ideally, if the training set contains the same type of occlusion as the test image then identification becomes a rather straightforward task. In practice, however, there is no guarantee that the collected data would cover all different occlusions for all identities of interest. An example of this problem is presented in Figure 1. The image database consists of non-occluded faces of subjects with intra-class illumination differences while the query face exhibits a 70% random block occlusion that covers most of the informative features of the face. In applications where prior knowledge such as the region and the object of occlusion is not provided, an appropriate modeling of the error between the test image and the training samples is necessary.

This paper has been partially supported by the Spanish Ministry of Economy and Competitiveness under project TIN2013-43880-R.

M. Iliadis and A. K. Katsaggelos are with the Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208-3118 USA (e-mail: miliad@northwestern.edu; aggk@eecs.northwestern.edu).

H. Wang is with TCL Research America, San Jose, CA 95131 USA (e-mail: haohong.wang@tcl.com).

R. Molina is with the Departamento de Ciencias de la Computación e I.A., Universidad de Granada, 18071 Granada, Spain (e-mail: rms@decsai.ugr.es).

Early works on face identification [2], [3] attempted to deal with illumination variations. To handle more complex variations such as face disguise and expressions, sparse representation-based classification models were proposed [4]–[8]. The main idea in these approaches is that a subject's face sample can be represented as a linear combination of available images of the same subject. Then, the face class that yields the minimum reconstruction error is selected in order to classify the subject.

To address cases with complex occlusion and corruption robust representation models [4], [9]–[13] of the error image[1] were considered, utilizing a non-Gaussian distribution model to minimize the influence of outliers. In these models a Laplacian distribution (sparse error) or more general distributions based on M-Estimators [11], [14] were fitted to the errors. There are, however, two main drawbacks with such approaches. First, the iterative reweighted algorithm used to solve the robust representation problem is computationally expensive when dealing with high-dimensional data [9]–[12]. Second, the performance degrades with over 50% block occlusion. According to [15], this is due to the assumption that error pixels are independent. The robust representation approaches are usually effective in FI cases with non-contiguous[2] errors such as pixel corruption.

In cases of contiguous errors there is a spatial correlation among the error pixels as mentioned in [15]–[17]. To exploit this correlation, the spatial continuity of the error image was integrated into the sparse representation-based classification model [15], [16]. However, these models lack convergence guarantees [17]. To address this issue Qian et. al. [17] observed that the error image with contiguous occlusion is low-rank and proposed to estimate the error support by solving a nuclear norm minimization problem with the use of ADMM [18]. While the low-rank assumption was well justified the method was effective with up to 50% random block occlusion. A reason might be that only the structure of the error (error support) was exploited and not its distribution (e.g., sparsity).

In this work, we propose an iterative method to solve the FI problem with occlusions. We consider the same scenario as in [9], [10], [17] and [15] according to which we are given frontal aligned views with a block occlusion which appear in any position on the test image but is "unseen" to the training data.

[1]Error image is the difference between the occluded test face and the unoccluded training face of the same identity.

[2]We call variations such as block occlusion and face disguise (e.g., scarves) contiguous errors since the error image is zero everywhere except in the region of the occluded object.

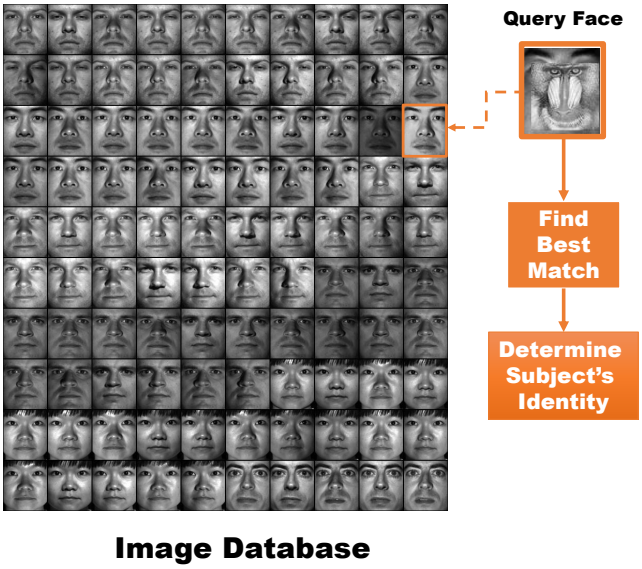

Fig. 1: Overview of the face identification problem investigated in this work. The image database consists of non-occluded faces of different subjects with small illumination changes while the query face exhibits 70% block occlusion that covers most of the informative features of the face.

As already mentioned, with the robust methods [9], [10], [15], [17] high computational cost is exhibited and identification results significantly degrade with over 50% block occlusion. Our approach is based on a new iterative method which is efficient in terms of computational cost and robust to block occlusions up to 70%. To describe contiguous errors (e.g., a random block occlusion) the proposed method utilizes a robust representation with two characteristics. The first fits to the errors a distribution based on a tailored loss function. The second models the low-rank structure of the error. We will then show that this joint characterization is effective for describing errors with spatial continuity. Our approach is efficient in terms of computational cost. Efficient minimization is performed by reformulating the reweighted coding problem as a constrained ADMM one, thus, avoiding costly matrix inversions.

A special case of our fast iterative algorithm leads to the robust representation method presented in [9], [10] which is normally used to handle non-contiguous errors (e.g., pixel corruption).

The rest of the paper is organized as follows. In Section II first we present our fast iterative algorithm for FI cases with occlusion. Then, we describe the identification scheme and the weight function used in our method. Finally, a special case of the algorithm for solving efficiently the FI problem with non-contiguous errors is described. Experimental results and discussion on the performance of the proposed algorithms are presented in Section III, and conclusions are drawn in Section IV.

II. Robust Face Identification with Block Occlusions

In this section we propose an iterative method to address the FI problem with block occlusions. Our approach utilizes the robust representation of [9]–[12] with two characteristics and uses a tailored loss function based on M-estimators. The method handles contiguous errors that are considered low-rank in comparison to the size of the image and is efficient in terms of computational cost. A special case of our method is also utilized to solve efficiently the robust representation problem for non-contiguous errors.

Let $\mathbf{y} \in \mathbb{R}^d$ denote the face test sample in a column-wise vectorized form where $j \times k = d$ is the size of the image. Let $\mathbf{T} = [T_i, \ldots, T_c] \in \mathbb{R}^{d \times n}$ denote the matrix (dictionary) with the set of samples of c subjects stacked in columns. $T_i \in \mathbb{R}^{d \times n_i}$ denotes the n_i set of samples of the i^{th} subject, such that, $\sum_i n_i = n$.

As illustrated in Figure 2 we can represent the test sample with occlusion as the superposition of training samples and a representation error $\mathbf{e}$, thus, the degradation model is,

$$\mathbf{y} = \mathbf{T}\mathbf{a} + \mathbf{e}, \tag{1}$$

where $\mathbf{e} \in \mathbb{R}^d$ is the representation error and $\mathbf{a} \in \mathbb{R}^n$ is the representation vector. Thus, the test sample can be represented as a linear combination of the samples in $\mathbf{T}$. The face identity is chosen based on the minimum class residual provided by the estimated coefficients $\mathbf{a}$ as in [4]. The residual image $\mathbf{e}$ in our model has two characteristics:

1) It includes mainly the occluded object, and therefore it is considered low-rank since many of its rows or columns are zero.
2) It follows a distribution that can be effectively described by a tailored potential loss function (e.g., the logistic function utilized in [10]).

We expect that the calculation of $\mathbf{e}$ based on the two characteristics mentioned above will lead to an accurate estimation of the $\mathbf{a}$ coefficients and provide the correct identity. Although the first characteristic has been employed in [17] and the second in robust representation methods such as [10], both of them are necessary to adequately describe the residual image and are used together for the first time in our work. In particular, enforcing only the second characteristic will not necessarily lead to an error image that is structured and low-rank. To the best of our knowledge there is no study that utilizes both the robust representation and the structure of the error using low-rank estimation in a unified framework.

For the problem described above, we propose the following function to be minimized,

$$J(\mathbf{a}) = \sum_{i=1}^{d} \phi((\mathbf{y} - \mathbf{T}\mathbf{a})_i) + \lambda_* \left\| T_M(\mathbf{y} - \mathbf{T}\mathbf{a}) \right\|_* + \vartheta(\mathbf{a}), \tag{2}$$

where $\phi : \mathbb{R} \to \mathbb{R}$ is a potential loss function which is selected from a variety of M-Estimators [11], [14] and $(\mathbf{y} - \mathbf{T}\mathbf{a})_i$ is the i^{th} component error. The function $\vartheta(\mathbf{a})$ is used as a regularizer of the coefficients $\mathbf{a}$, $\lambda_* > 0$ and T_M is an operator that transforms its vector argument into a matrix of appropriate size such that $T_M(\mathbf{y} - \mathbf{T}\mathbf{a}) \in \mathbb{R}^{j \times k}$. The nuclear norm $\|\cdot\|_*$ is the convex-relaxation of the rank function.

Thus we are looking in minimizing the following problem,

$$\min_{\mathbf{a}} J(\mathbf{a}). \tag{3}$$

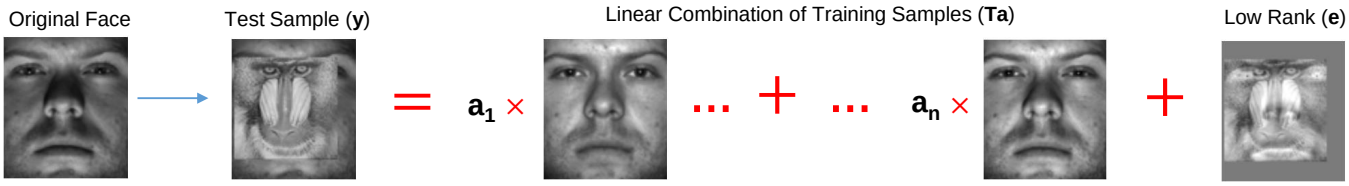

Fig. 2: Degradation Model: the test sample with occlusion can be represented as the linear combination of training samples with some intra-class variations (e.g., lighting) plus the error $\mathbf{e}$. The error $\mathbf{e}$ has two characteristics; it is considered low-rank and follows a distribution described by a tailored potential loss function.

Notice, how previous works relate to our model:

1) For $\phi(x) = x^2$, $\lambda_* = 0$ and $\vartheta(\mathbf{a}) = \lambda||\mathbf{a}||_1$ with $\lambda > 0$, it is the sparse representation-based classification (SRC) [4] given as,

$$\min_{\mathbf{a}} \|\mathbf{y} - \mathbf{T}\mathbf{a}\|_2^2 + \lambda||\mathbf{a}||_1. \tag{4}$$

2) For $\phi(x) = x^2$ and $\vartheta(\mathbf{a}) = \lambda||\mathbf{a}||_2^2$, it is the low-rank regularized regression (LR^3) [17] which is formulated as,

$$\min_{\mathbf{a}} \|\mathbf{y} - \mathbf{T}\mathbf{a}\|_2^2 + \lambda_* \|T_M(\mathbf{y} - \mathbf{T}\mathbf{a})\|_* + \lambda \|\mathbf{a}\|_2^2. \tag{5}$$

3) For $\lambda_* = 0$, it is the robust representation problem [9]–[12] formulated as,

$$\min_{\mathbf{a}} \sum_{i=1}^{d} \phi((\mathbf{y} - \mathbf{T}\mathbf{a})_i) + \vartheta(\mathbf{a}). \tag{6}$$

In previous works authors have chosen different functions $\vartheta(\mathbf{a})$ to regularize the coefficients $\mathbf{a}$. In the collaborative representation-based classification with regularized least square (CR-RLS) [7] the authors are solving the SRC problem with $\vartheta(\mathbf{a}) = \lambda||\mathbf{a}||_2^2$. In [9], [11] $\vartheta(\mathbf{a}) = \lambda||\mathbf{a}||_1$ was used combined with different potential functions while in the regularized robust coding (RRC) [10], $\vartheta(\mathbf{a}) = \lambda||\mathbf{a}||_2^2$ was used. In correntropy-based sparse representation (CESR) [12] and structured sparse error coding (SSEC) [15], $\vartheta(\mathbf{a})$ was chosen to be the indicator function of the nonnegative orthant $\mathbb{R}_+^n$, such that a nonnegative $\mathbf{a} \geq 0$ regularization term was enforced.

In this work we choose $\vartheta(\mathbf{a})$ so as the representation coefficients are nonnegative, since it has been shown to provide robust representation for FI in [12] and [15].

The robustness property of $\phi(\cdot)$ in (2), to be described now, in combination with $\vartheta(\cdot)$ will force some of the coefficients of $\mathbf{e}$ to be zero or very small in magnitude. It will also force $\mathbf{a}$ to be concentrated in areas of the images that can be represented well by faces in $\mathbf{T}$. Furthermore, the use of the nuclear norm will force the residual to be low-rank.

We consider potential loss functions $\phi(\cdot)$ symmetric around zero associated to Super Gaussian (SG) distributions [19]. The function $\phi(\sqrt{x})$ has to be increasing and concave for $x \in (0, \infty)$ [20]. This condition is equivalent to $\phi'(x)/x$ being decreasing on $(0, \infty)$, that is, for $x_1 \geq x_2 \geq 0$, $\phi'(x_1)/x_1 \leq \phi'(x_2)/x_2$. If this condition is satisfied, then $\phi(\cdot)$ can be represented as (using [21, Ch. 12]),

$$\phi(x) = \inf_{\xi > 0} \frac{1}{2}\xi x^2 - \phi^*\left(\frac{1}{2}\xi\right), \tag{7}$$

where $\phi^*(\xi)$ is the concave conjugate of $\phi(\sqrt{x})$ and ξ is a variational parameters. The dual relationship to (7) is given by [21],

$$\phi^*\left(\frac{1}{2}\xi\right) = \inf_{x} \frac{1}{2}\xi x^2 - \phi(x). \tag{8}$$

Equality in (7) is obtained at the optimal values of ξ, which are computed from the dual representation (8) by taking the derivative with respect to x and setting it to zero, which gives $\xi = \phi'(x)/x$. By using (7) we can write the function in (2) as,

$$\begin{aligned} J(\mathbf{a}) = &\min_{\mathbf{w}} \frac{1}{2}\|\sqrt{\mathbf{W}}(\mathbf{y} - \mathbf{T}\mathbf{a})\|_2^2 + \varphi(\mathbf{w}) \\ &+ \lambda_* \|T_M(\mathbf{y} - \mathbf{T}\mathbf{a})\|_* + \vartheta(\mathbf{a}), \end{aligned} \tag{9}$$

where $\mathbf{w} = (\xi_1, \ldots, \xi_d)$ with $\xi_i > 0, i = 1, \ldots, d$, $\mathbf{W} = \text{diag}(\mathbf{w})$, and $\varphi(\mathbf{w}) = \sum_{i=1}^{d} \phi^*(\frac{1}{2}\xi_i)$.

Notice, before proceeding, that the weights in $\mathbf{W}$ are of the form $\phi'(x)/x$ which are large for small values of x for SG potential loss functions, so $\mathbf{a}$ will fit well small values of $\mathbf{e}$ in magnitude.

Let us consider the augmented function by subsistuting (9) into (3),

$$\begin{aligned} J(\mathbf{a}, \mathbf{w}) = &\frac{1}{2}\|\sqrt{\mathbf{W}}(\mathbf{y} - \mathbf{T}\mathbf{a})\|_2^2 + \varphi(\mathbf{w}) \\ &+ \lambda_* \|T_M(\mathbf{y} - \mathbf{T}\mathbf{a})\|_* + \vartheta(\mathbf{a}). \end{aligned} \tag{10}$$

A local minimizer $(\mathbf{a}, \mathbf{w})$ can be calculated by alternating minimization in two steps; in step one, the weights are updated by fixing the representation coefficients $\mathbf{a}$ and in step two the vector $\mathbf{a}$ is updated by fixing the weights in $\mathbf{W}$, i.e.,

$$w_i^{t+1} = \phi'((\mathbf{y} - \mathbf{T}\mathbf{a}^t)_i)/(\mathbf{y} - \mathbf{T}\mathbf{a}^t)_i \tag{11a}$$

$$\begin{aligned} \mathbf{a}^{t+1} = &\operatorname*{argmin}_{\mathbf{a}} \|\sqrt{\mathbf{W}^{t+1}}(\mathbf{y} - \mathbf{T}\mathbf{a})\|_2^2 \\ &+ \lambda_* \|T_M(\mathbf{y} - \mathbf{T}\mathbf{a})\|_* + \vartheta(\mathbf{a}), \end{aligned} \tag{11b}$$

where w_i^{t+1} and $\mathbf{a}^{t+1}$ are the weights and representation coefficients estimated at the t^{th} iteration, respectively. The term $(\mathbf{y} - \mathbf{T}\mathbf{a}^t)_i$ denotes the component error i at the t^{th} iteration. Large weights are assigned to pixels in the residual image with small errors in the previous reconstructed iteration, while small weights are assigned to pixels in the residual image with large errors.

We expect the use of the nuclear norm of the residual to force small weights in $\mathbf{W}$ to be assigned only on the occluded part. Thus, outlier pixels will not contribute much to the

Fig. 3: An example of 70% occluded face. If the low-rank constraint is not used then $\mathbf{y}$ is reconstructed from faces that belong to the wrong identities. On the other hand, our method is able to reconstruct the occluded face from images that belong to the correct identity.

reconstruction at the next iteration, and at convergence, the estimated error will mainly consist of those outliers.

Notice that we have,

$$J(\mathbf{a}^t) = J(\mathbf{a}^t, \mathbf{w}^{t+1}) \geq J(\mathbf{a}^{t+1}, \mathbf{w}^{t+1}) \geq J(\mathbf{a}^{t+1}, \mathbf{w}^{t+2}) = J(\mathbf{a}^{t+1}). \quad (12)$$

A. Optimization

Let us now describe the iterative algorithm to find efficiently $\mathbf{a}^{t+1}$ in problem (11b).

In order to solve the proposed problem, first we let $\mathbf{y} - \mathbf{T}\mathbf{a} = \mathbf{e}$ and since we are interested in estimating non-negative coefficients for the representation vector we also introduce an additional variable $\mathbf{z}$ such that $\mathbf{a} = \mathbf{z}$. Then, the coding step (11b) is reformulated as,

$$\begin{aligned} &\underset{\mathbf{a},\mathbf{z},\mathbf{e}}{\text{minimize}} \quad \|\sqrt{\mathbf{W}^{t+1}}\mathbf{e}\|_2^2 + \lambda_* \|T_M(\mathbf{e})\|_* + \vartheta(\mathbf{z}) \\ &\text{subject to} \quad \mathbf{y} - \mathbf{T}\mathbf{a} = \mathbf{e}, \mathbf{a} = \mathbf{z}. \end{aligned} \quad (13)$$

Problem (13) is solved efficiently with ADMM which is known for fast convergence to an approximate solution [18]. As in the method of multipliers, the problem takes the form of the augmented Lagrangian,

$$\begin{aligned} L(\mathbf{e}, \mathbf{a}, \mathbf{z}, \mathbf{u}_1, \mathbf{u}_2, \mathbf{w}^{t+1}) &= \|\sqrt{\mathbf{W}^{t+1}}\mathbf{e}\|_2^2 + \lambda_* \|T_M(\mathbf{e})\|_* \\ &+ \vartheta(\mathbf{z}) + \mathbf{u}_1^T\big(\mathbf{y} - \mathbf{T}\mathbf{a} - \mathbf{e}\big) + \frac{\rho_1}{2}\|\mathbf{y} - \mathbf{T}\mathbf{a} - \mathbf{e}\|_2^2 \\ &+ \mathbf{u}_2^T\big(\mathbf{a} - \mathbf{z}\big) + \frac{\rho_2}{2}\|\mathbf{a} - \mathbf{z}\|_2^2, \end{aligned} \quad (14)$$

where $\rho_1 > 0$ and $\rho_2 > 0$ are the penalty parameters, and $\mathbf{u}_1$ and $\mathbf{u}_2$ are the dual variables. The ADMM updates can be expressed as,

$$\mathbf{e}_{s+1} = \underset{\mathbf{e}}{\operatorname{argmin}}\, L(\mathbf{e}, \mathbf{a}_s, \mathbf{u}_{1,s}, \mathbf{w}^{t+1}), \quad (15a)$$

$$\mathbf{z}_{s+1} = \underset{\mathbf{z}}{\operatorname{argmin}}\, L(\mathbf{z}, \mathbf{a}_s, \mathbf{u}_{2,s}), \quad (15b)$$

$$\mathbf{a}_{s+1} = \underset{\mathbf{a}}{\operatorname{argmin}}\, L(\mathbf{e}_{s+1}, \mathbf{a}, \mathbf{u}_{1,s}, \mathbf{z}_{s+1}, \mathbf{u}_{2,s}), \quad (15c)$$

$$\mathbf{u}_{1,s+1} = \mathbf{u}_{1,s} + \rho_1(\mathbf{y} - \mathbf{T}\mathbf{a}_{s+1} - \mathbf{e}_{s+1}), \quad (15d)$$

$$\mathbf{u}_{2,s+1} = \mathbf{u}_{2,s} + \rho_2(\mathbf{a}_{s+1} - \mathbf{z}_{s+1}), \quad (15e)$$

where s denotes the ADMM iteration and finally we set $\mathbf{a}^{t+1} = \lim_{s\to+\infty} \mathbf{a}_{s+1}$.

1) Finding $\mathbf{e}_{s+1}$: The update of $\mathbf{e}_{s+1}$ is given by minimizing the following problem,

$$\begin{aligned} \mathbf{e}_{s+1} = \underset{\mathbf{e}}{\operatorname{argmin}} \frac{1}{2}\|\mathbf{y} - \mathbf{T}\mathbf{a}_s - \mathbf{e}\|_2^2 + \frac{1}{\rho_1}\|\sqrt{\mathbf{W}^{t+1}}\mathbf{e}\|_2^2 \\ + \frac{1}{\rho_1}\mathbf{u}_{1,s}^T\big(\mathbf{y} - \mathbf{T}\mathbf{a}_s - \mathbf{e}\big) + \frac{\lambda_*}{\rho_1}\|T_M(\mathbf{e})\|_*. \end{aligned} \quad (16)$$

To calculate $\mathbf{e}_{s+1}$ we consider a two-step fast approximation. In step one we solve the weighted norm problem,

$$\begin{aligned} \tilde{\mathbf{e}} = \underset{\mathbf{e}}{\operatorname{argmin}} \frac{1}{2}\|\mathbf{y} - \mathbf{T}\mathbf{a}_s - \mathbf{e}\|_2^2 + \frac{1}{\rho_1}\|\sqrt{\mathbf{W}^{t+1}}\mathbf{e}\|_2^2 \\ + \frac{1}{\rho_1}\mathbf{u}_{1,s}^T\big(\mathbf{y} - \mathbf{T}\mathbf{a}_s - \mathbf{e}\big), \end{aligned} \quad (17)$$

and in step two to satisfy the nuclear norm constraint we project the estimated $\tilde{\mathbf{e}}$ to a low-rank space, to obtain,

$$\mathbf{e}_{s+1} = \underset{\mathbf{e}}{\operatorname{argmin}} \frac{1}{2}\Big\|T_M\Big(\mathbf{e} - \tilde{\mathbf{e}}\Big)\Big\|_F^2 + \frac{\lambda_*}{\rho_1}\|T_M(\mathbf{e})\|_*. \quad (18)$$

Problem (17) has a closed-form solution given by,

$$\tilde{\mathbf{e}} = \mathbf{C}^{-1}\big(\mathbf{y} - \mathbf{T}\mathbf{a}_s + \mathbf{u}_{1,s}/\rho_1\big), \quad (19)$$

where $\mathbf{C} = \big(\mathbf{I} + 2\mathbf{W}^{t+1}/\rho_1\big)$ is a diagonal matrix with diagonal entries $c_i = 1 + 2w_i^{t+1}/\rho_1$. Since $\mathbf{C}$ is a diagonal matrix, to update $\tilde{\mathbf{e}}$ we only need to construct a vector with elements equal to $1/c_i$ and perform an element wise-multiplication between the constructed vector and the residual vector $\mathbf{y} - \mathbf{T}\mathbf{a}_s + \mathbf{u}_{1,s}/\rho_1$. Thus, the update of $\tilde{\mathbf{e}}$ can be calculated fast.

The update in (19) is in essence an outlier detector similar to the *soft-thresholding* operator [18]. The values of the residual vector $\mathbf{y} - \mathbf{T}\mathbf{a}_s + \mathbf{u}_{1,s}/\rho_1$ will be weighted according to $\mathbf{C}^{-1}$. A small weight (close to zero) will be given to non-outlier elements (e.g., elements of the residual vector with small values) while a large weight (close to one) will be given to outliers (e.g., elements of the residual vector with large values).

Problem (18) is a nuclear norm minimization problem of the form,

$$\min_{\mathbf{X}} \frac{1}{2}\|\mathbf{X} - \Delta\|_F^2 + \lambda_*\|\mathbf{X}\|_*, \quad (20)$$

which has a closed-form solution given by [18],

$$\hat{\mathbf{X}} = \mathcal{L}_{\lambda_*}(\Delta) = \mathbf{U}_\Delta \mathcal{S}_{\lambda_*}(\mathbf{\Sigma}_\Delta)\mathbf{V}_\Delta^T, \tag{21}$$

where $\mathcal{S}_{\lambda_*}(\mathbf{\Sigma}_\Delta) = \text{sign}(\delta_{ij})\max(0, |\delta_{ij}| - \lambda_*)$ is the soft-thresholding operator, $\mathcal{L}_{\lambda_*}(\cdot)$ is the singular value soft-thresholding operator and $\Delta = \mathbf{U}_\Delta \mathbf{\Sigma}_\Delta \mathbf{V}_\Delta^T$ is the SVD of matrix Δ.

Thus, the two-step solution of (16) is given by,

$$\tilde{\mathbf{e}} = \mathbf{C}^{-1}\big(\mathbf{y} - \mathbf{T}\mathbf{a}_s + \mathbf{u}_{1,s}/\rho_1\big) \tag{22a}$$

$$\mathbf{e}_{s+1} = \mathcal{L}_{\lambda_*/\rho_1}\big(T_M(\tilde{\mathbf{e}})\big). \tag{22b}$$

The solution in (22) is the low-rank estimation of the weighted error image.

2) Finding $\mathbf{z}_{s+1}$: The update of $\mathbf{z}_{s+1}$ is obtained by solving the following problem,

$$\mathbf{z}_{s+1} = \underset{\mathbf{z}}{\text{argmin}} \frac{1}{2}\|\mathbf{a}_s - \mathbf{z}\|_2^2 + \frac{1}{\rho_2}\vartheta(\mathbf{z}) + \frac{1}{\rho_2}\mathbf{u}_{2,s}^T\big(\mathbf{a}_s - \mathbf{z}\big). \tag{23}$$

The solution of (23) is given by,

$$\mathbf{z}_{s+1} = (\mathbf{a}_s + \mathbf{u}_{2,s}/\rho_2)_+, \tag{24}$$

where $(\cdot)_+$ is the function the keeps only the positive coefficients of its argument and set the rest to zero.

Notice that we only need to change the update of $\mathbf{z}$ for different regularization functions $\vartheta(\mathbf{z})$. For example, to solve an iterative reweighted sparse coding (IRSC) problem and regularize the coefficients to be sparse ($\vartheta(\mathbf{z}) = \lambda||\mathbf{z}||_1$) we have to substitute (24) with the *soft-thresholding* operator [18]. To solve an iterative reweighted least squares (IRLS) problem ($\vartheta(\mathbf{a}) = \lambda||\mathbf{a}||_2^2$) we do not need to introduce and estimate the additional variables, $\mathbf{z}$ and $\mathbf{u}_2$. The coefficients $\mathbf{a}$ can be estimated by solving a regularized least squares problem.

3) Finding $\mathbf{a}_{s+1}$: The update of $\mathbf{a}_{s+1}$ is obtained by solving the following problem,

$$\begin{aligned}\mathbf{a}_{s+1} = \underset{\mathbf{a}}{\text{argmin}} \frac{1}{2}\|\mathbf{y} - \mathbf{T}\mathbf{a} - \mathbf{e}_{s+1}\|_2^2 + \frac{1}{\rho_1}\mathbf{u}_{1,s}^T\big(\mathbf{y} - \mathbf{T}\mathbf{a} - \mathbf{e}_{s+1}\big) \\ + \frac{1}{\rho_1}\mathbf{u}_{2,s}^T\big(\mathbf{a} - \mathbf{z}_{s+1}\big) + \frac{\rho_2}{2\rho_1}\|\mathbf{a} - \mathbf{z}_{s+1}\|_2^2.\end{aligned} \tag{25}$$

Equation (25) has a closed-form solution,

$$\mathbf{a}_{s+1} = \mathbf{P}\Big(\mathbf{T}^T(\mathbf{y} - \mathbf{e}_{s+1} + \mathbf{u}_{1,s}/\rho_1) + (\rho_2/\rho_1)\mathbf{z} - \mathbf{u}_{2,s}/\rho_1\Big), \tag{26}$$

where $\mathbf{P} = \Big(\mathbf{T}^T\mathbf{T} + \frac{\rho_2}{\rho_1}\mathbf{I}\Big)^{-1}$ and can be pre-calculated once and cached offline.

An example of our approach is presented in Figure 3. If the low-rank constraint is not used then $\mathbf{y}$ is reconstructed using faces that belong to the wrong identities. On the other hand, our approach with the use of the nuclear norm constraint is able to reconstruct the occluded face using images that belong to the correct identity (second row in Figure 3). Also, notice the differences in weight map estimations. In the first case small weights (black values) assigned without any structure to any region of the face. This is not desirable since informative pixels were detected as outliers (e.g., pixels around the occluded object). In addition, many pixels on the occluded object were detected as inliers. In the second case, the error is low-rank and has a spatial continuity around the occluded object (pixels close to zero). In this case, the weight map $\mathbf{W}$ is also enforced to have a spatial continuity (since $\mathbf{W}$ is related to the error) with small weights assigned to the occluded object as desired.

We would like to mention here a closely related work for block occlusion errors namely robust low-rank regularized regression (RLR^3) [17]. There are two fundamental differences between our work and RLR^3. First, the weighted residual $\mathbf{W}(\mathbf{y} - \mathbf{T}\mathbf{a})$ instead of the residual error $\mathbf{y} - \mathbf{T}\mathbf{a}$ was modeled to be low-rank which is different from our method. The model in RLR^3 can handle occlusion of specific objects and size that covers a portion of the face image entirely from left to right (or from top to bottom) such as scarves. Our method handles block occlusions that appear in any size and place in the face. Second, the function to be minimized in RLR^3 is not derived from the duality theorem [21] which raises concerns about its convergence guarantees.

The complete steps of the fast low-rank and iterative reweighted nonnegative least squares (F-LR-IRNNLS) algorithm for contiguous errors are presented in Algorithm 1.

Algorithm 1 Fast & Low-Rank IRNNLS Algorithm

Inputs: $\mathbf{y}$, $\mathbf{T}$, λ_*, ρ_1, ρ_2, ϵ_1, ϵ_2 and ϵ_3.
Initialize $\mathbf{a}^1 = 1/n$, $\mathbf{u}_{1,1} = \mathbf{0}$, $\mathbf{u}_{2,1} = \mathbf{0}$, $t = 0$ and $s = 0$
Repeat
1) $t = t + 1$
2) Estimate the weights,

$$w_i^t = \phi'((\mathbf{y} - \mathbf{T}\mathbf{a}^t)_i)/(\mathbf{y} - \mathbf{T}\mathbf{a}^t)_i, \;\; i = 1, \ldots, d$$

Repeat
3) $s = s + 1$
4) Find $\mathbf{e}_s$ using (22) (contiguous errors) or (32) (non-contiguous errors)
5) Find $\mathbf{z}_s$ using (24)
6) Update $\mathbf{a}_s$ using (26)
7) Update $\mathbf{u}_{1,s}$ and $\mathbf{u}_{2,s}$ using (15d), (15e)

Until *converge*
8) Set $\mathbf{a}^t = \mathbf{a}_s$, $\mathbf{u}_{1,1} = \mathbf{0}$, $\mathbf{u}_{2,1} = \mathbf{0}$, $s = 0$
Until *converge*
Output: The final estimates of $\mathbf{a}$ and $\mathbf{w}$.

B. Identification Scheme

In SRC [4], the face class that yields the minimum reconstruction error is selected in order to classify or identify the subject. Similarly, in this work the classification is given by computing the residuals e for each class i as,

$$e_i(\mathbf{y}) = \|\sqrt{\mathbf{W}^f}(\mathbf{y} - T_i\hat{\mathbf{a}}_i)\|_2, \tag{27}$$

where $\hat{\mathbf{a}}_i$ is the segment of the final estimated $\mathbf{a}$ associated with class i and $\mathbf{W}^f$ is the final estimated weight matrix from Algorithm 1. Finally, the identity of $\mathbf{y}$ is given as, $\text{Identity}(\mathbf{y}) = \text{argmin}_i\{e_i(\mathbf{y})\}$.

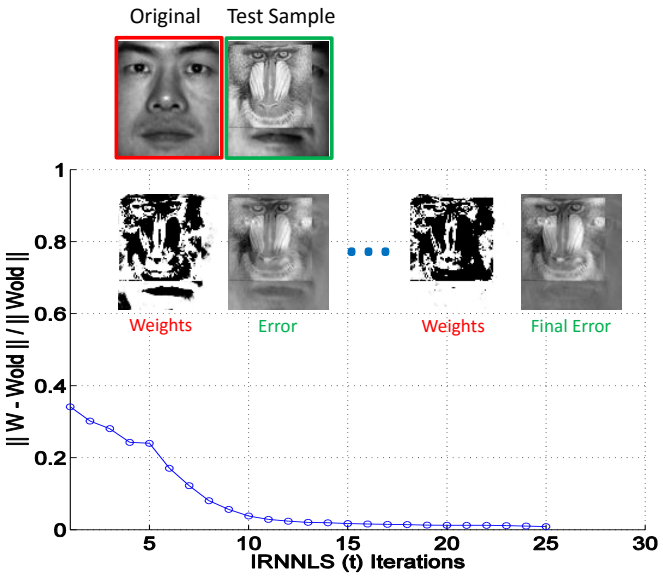

Fig. 4: The graph shows the convergence of the iterative reweighted and low-rank algorithm (F-LR-IRNNLS). In particular, we present the change of the weights between two consecutive iterations. As $t \to \infty$ the error becomes more sparse and structural, thus, small weights are concentrated on the occluding object.

C. The Weight Function

Ideally, the weight function should distinguish inliers and outliers given a training dictionary with non-occluded faces and a test sample with occlusion [10]. In particular, given the residual error at any iteration, small weights (close to zero) should be assigned to the outlier pixels (large residual error) and larger weight (close to one) to the inlier pixels (small residual error). Although any weight function [11], [14] of the form $w = \phi'(x)/x$ can be used in our framework as long as $\phi'(x)/x$ is decreasing on $(0, \infty)$, in this work we utilize the logistic function proposed in [10]. The logistic function performs particularly well in FI[3]. The weight component w_i as a function of x_i, which is decreasing on $(0, \infty)$, is given by,

$$w_i = \frac{\exp\left(-\mu\left(x_i\right)^2 + \mu\eta\right)}{1 + \exp\left(-\mu\left(x_i\right)^2 + \mu\eta\right)}, \quad i = 1, \ldots, d, \tag{28}$$

where μ and η are positive scalars. As in [10], η denotes the value of the l^{th} largest element of the residual vector $\mathbf{x}$, where $l = \lfloor \gamma d \rfloor$, $\gamma \in (0, 1)$. μ is given as $\frac{\zeta}{\eta}$ with $\zeta = 8$. We also set $\gamma = 0.8$ for the experiments without occlusion and $\gamma = 0.6$ for the experiments with occlusion as in [10].

D. Convergence Criteria

For the purpose of this paper, in order to guarantee convergence of the optimization problem (13) using ADMM, it is sufficient to enforce appropriate termination criteria. As suggested in [18], we enforced the primal residuals to be small such that $\|\mathbf{y} - \mathbf{T}\mathbf{a} - \mathbf{e}\|_2 \leq \epsilon_1$ and $\|\mathbf{a} - \mathbf{z}\|_2 \leq \epsilon_2$, where ϵ_1 and ϵ_2 are small positive numbers.

The termination criterion for the iterative reweighted sequences is $\left\|\mathbf{W}^t - \mathbf{W}^{t-1}\right\|_2 / \left\|\mathbf{W}^{t-1}\right\|_2 < \epsilon_3$, where ϵ_3 is a small positive number. Figure 4 shows the change of the weights between two consecutive iterations for one particular example.

E. Special case: robust representation for non-contiguous errors

A special case of our method leads to the robust representation method [9]–[12] when used for FI problem with non-contiguous errors as given in (6). In the previous robust methods [9]–[12] high computational cost is exhibited due to the reweighted coding step to estimate $\mathbf{a}$. To address this issue the proposed ADMM algorithm described above is utilized here as well to solve this case efficiently. The augmented function of the robust representation problem is formulated as,

$$J(\mathbf{a}, \mathbf{w}) = \frac{1}{2}\|\sqrt{\mathbf{W}}(\mathbf{y} - \mathbf{T}\mathbf{a})\|_2^2 + \varphi(\mathbf{w}) + \vartheta(\mathbf{a}), \tag{29}$$

which is similar to (10) but without the nuclear norm term. Similarly to (10), a local minimizer $(\mathbf{a}, \mathbf{w})$ can be calculated in two steps,

$$w_i^{t+1} = \phi'((\mathbf{y} - \mathbf{T}\mathbf{a}^t)_i)/(\mathbf{y} - \mathbf{T}\mathbf{a}^t)_i \tag{30a}$$

$$\hat{\mathbf{a}}^{t+1} = \underset{\mathbf{a}}{\operatorname{argmin}} \|\sqrt{\mathbf{W}^{t+1}}(\mathbf{y} - \mathbf{T}\mathbf{a})\|_2^2 + \vartheta(\mathbf{a}). \tag{30b}$$

A major drawback of the iterative reweighted algorithm, especially in large-scale FI problems, is that it is computationally expensive [22] due to the coding step (30b). High computational cost is exhibited regardless of the coefficient regularization, for example $\vartheta(\mathbf{a}) = \lambda||\mathbf{a}||_2^2$, $\vartheta(\mathbf{a}) = \lambda||\mathbf{a}||_1$ or $\vartheta(\mathbf{a})$ is the indicator function of the nonnegative orthant $\mathbb{R}_+^n$ [9]–[12]. The coding step is expensive since in each iteration a new weighted system matrix $\sqrt{\mathbf{W}^{t+1}}\mathbf{T}$ is provided given the updated weights. Moreover, an offline pre-calculation of the weighted inverse is not possible. Efficient methods to solve problems in the form of (30b) such as conjugate gradient [23] or ℓ_1 algorithms [24], [25] may require several iterations to converge to the desired point [10]. The active set method [26] used for solving the nonnegative least squares problem becomes also slow due to the computation of the pseudoinverse of the system matrix in each iteration.

To solve (30b) efficiently for the $\vartheta(\mathbf{a})$ functions described earlier we utilize the proposed ADMM algorithm described above. Thus, we avoid the explicit calculation of $\mathbf{T}^T\mathbf{W}^{t+1}\mathbf{T}$ inverse during the execution of the algorithm. The idea of accelerating the iterative reweighed scheme using ADMM is also found in [22]. However, in [22] the weights applied to the regularization term $\vartheta(\mathbf{a})$ rather to the residual as in done here. To accelerate (30b), we set $\mathbf{y} - \mathbf{T}\mathbf{a} = \mathbf{e}$, $\mathbf{a} = \mathbf{z}$ and the problem is reformulated as,

$$\begin{aligned} &\underset{\mathbf{a},\mathbf{z},\mathbf{e}}{\text{minimize}} \quad \|\sqrt{\mathbf{W}^{t+1}}\mathbf{e}\|_2^2 + \vartheta(\mathbf{z}) \\ &\text{subject to} \quad \mathbf{y} - \mathbf{T}\mathbf{a} = \mathbf{e}, \mathbf{a} = \mathbf{z}. \end{aligned} \tag{31}$$

[3]For a complete justification of the effectiveness of the logistic function in FI we refer the reader to [10].

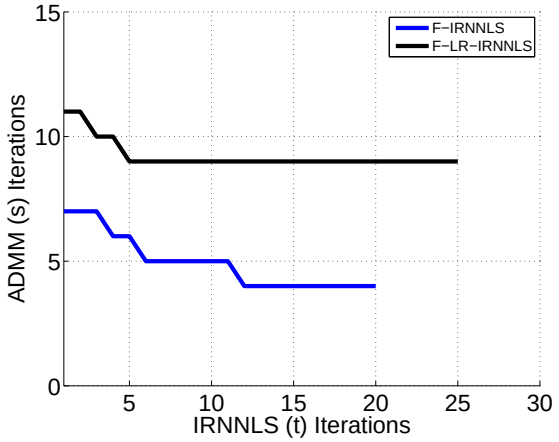

Fig. 5: The graph shows the number of ADMM iterations as a function of the reweighted iterations.

Notice that the $\sqrt{\mathbf{W}^{t+1}}\mathbf{T}$ term is no longer part of the optimization problem which allows us to solve (31) efficiently.

Problem (31) has similar ADMM updates with (13) except $\mathbf{e}_{s+1}$. To find $\mathbf{e}_{s+1}$ we only have to calculate,

$$\mathbf{e}_{s+1} = \mathbf{C}^{-1}\big(\mathbf{y} - \mathbf{T}\mathbf{a}_s + \mathbf{u}_{1,s}/\rho_1\big), \tag{32}$$

where $\mathbf{C}$ is the diagonal matrix defined in (19) and $\mathbf{C}^{-1}$ can be calculated fast due to the diagonal structure. The updates of $\mathbf{z}_{s+1}$ and $\mathbf{a}_{s+1}$ are similar to (24) and (26) respectively and as explained earlier $\mathbf{P}$ matrix can be pre-calculated once and cached offline. Thus, to solve the robust representation problem we utilize an efficient method since no online inversion of the system matrix is performed in any of the variable updates.

Our fast iterative reweighted nonnegative least squares (F-IRNNLS) algorithm just described solves the robust representation problem efficiently and the steps are also presented in Algorithm 1.

The number of ADMM iterations required for each reweighted iteration is presented in Figure 5 for our method. As expected the number of ADMM iterations for F-LR-IRNNLS is greater than the F-IRNNLS due to the calculation of the nuclear term.

III. Experimental Results

In this section we present experiments on three publicly available databases, AR [27], Extended Yale B [28] and Multi-PIE [29], to show the efficacy of the proposed method. We demonstrate identification and reconstruction results under various artificial and real-world variations. We compare our framework with eight other FI algorithms, SRC [4], CR-RLS [7], LR^3 [17], and the robust algorithms HQ (additive form) [11], CESR [12], RRC_L1 [10], RRC_L2 [10], SSEC [15]. We consider the following three FI cases:

1) cases with contiguous variations such as random block occlusion with different sizes and objects, face disguise and mixture noise which is a combination of block occlusion and pixel corruption,
2) cases with illumination variations,
3) cases with non-contiguous variations such as pixel corruption and face expressions.

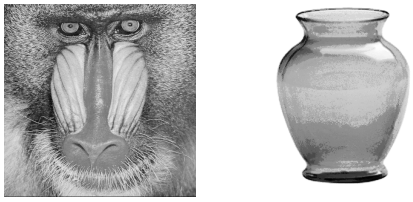

(a) Baboon image (b) Non-square image

Fig. 6: The two artificial images used for the block occlusion experiments.

For all methods, we used the solvers provided by the authors of the corresponding papers. We chose to solve the ℓ_1 minimization problem in SRC and RRC_L1 with the Homotopy algorithm[4] [30] since it resulted in the highest accuracy in the performance comparison in [24] with reasonable time execution. In our algorithms, we set $\lambda_* = 0.05$, $\rho_1 = 1$ and $\rho_2 = 0.1$. The convergence parameters were set equal to $\epsilon_1 = 10^{-2}$, $\epsilon_2 = 10^{-1}$, $\epsilon_3 = 10^{-2}$. For fair comparisons with respect to execution time and identification rates we set the same ϵ_3 for the RRC algorithm and the same maximum number of iterations ($t = 100$). All face images were normalized to have unit ℓ_2-norm and all variables initialized to zero except for $\mathbf{a}^1 = 1/n$ as in [10].

A. Identification under Block Occlusions

Experiments with occluded images were conducted on two datasets: Extended Yale B and AR.

As in [4], [10], [17], we chose Subsets 1 and 2 of Extended Yale B for training and Subset 3 for testing. Images were resized to 96×84 pixels. We considered two different artificial objects to occlude the test images as shown in Figure 6. For the first object, block occlusion was tested by placing the square baboon image on each test image. The location of the occlusion was randomly chosen and was unknown during training. We considered different sizes of the object such that the face is covered with the occluded object from 30% to 90% of its area. Identification rates for the different levels of occlusion are shown in Figure 7(a). For the second non-square and smooth (e.g., without textures in it) object shown in Figure 6(b), block occlusion was tested by randomly placing the object on each test image. Identification rates are shown in Figure 7(b).

To examine the robustness of our method across multiple datasets we evaluate its performance to block occlusion in the AR database. We chose the 7 non-occluded AR images per subject (each one with a different face expression) from session 1 for training and the 7 non-occluded images per subject from session 2 for testing. This experiment is more challenging since training faces appear with different expressions. In each test image, we replace a random block with the square baboon image. We resized the images to 60×43 pixels. The occlusion ratio increases from 30% to 50% and identification rates are shown in Figure 7(c).

[4]The source code of Homotopy algorithm can be downloaded at http://www.eecs.berkeley.edu/ yang/software/l1benchmark/

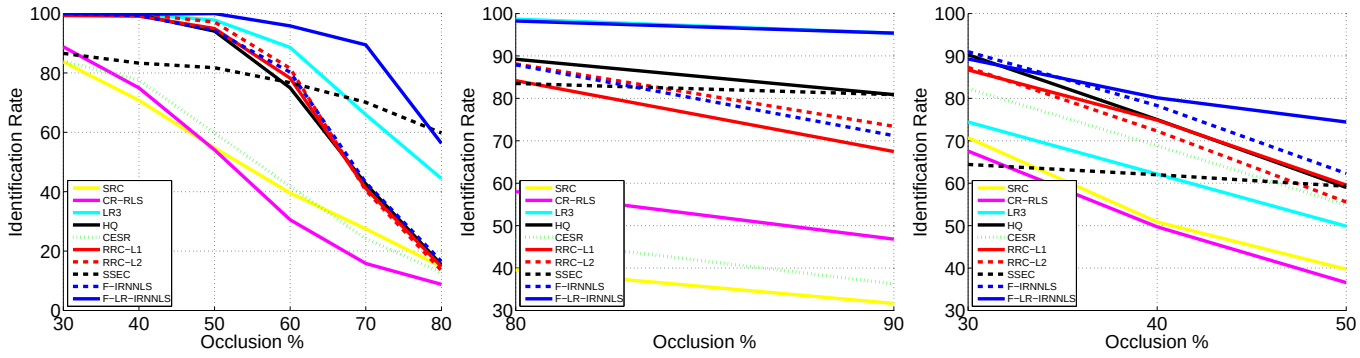

(a) Yale B dataset with baboon image (b) Yale B dataset with vase image (c) AR dataset with baboon image

Fig. 7: Identification Rates on the Extended Yale B and AR datasets under Block Occlusion with Baboon and Vase.

From the results with the baboon image we conclude that methods robust to contiguous errors (LR^3, SSEC) performed better than the non-contiguous methods in Yale B. Our F-LR-IRNNLS algorithm outperformed all previous methods overall in both Yale B and AR datasets. SSEC performed well with 80% occlusion in Yale B but it performed poorly in AR and with lower levels of occlusion. As explained in [17], there are no convergence guarantees for SSEC and, perhaps, this explains the unstable results obtained by this method in our experiments. Finally, the performance of non-contiguous error methods, HQ, RRCs and F-IRNNLS significantly dropped by high levels of occlusion. This is due to the fact that these methods cannot handle contiguous variations effectively.

From the results with the non-square object we found out that all methods performed better than when the baboon was used. We attribute this to the fact that the baboon object exhibits a lot of textures and that it looks like a face. Thus, it is much more difficult for the methods to distinguish inliers and outliers. LR^3 and F-LR-IRNNLS achieved similar identification rates at all levels of occlusion. Perhaps, modeling the error to be low-rank was sufficient for the non-square case in the Yale B dataset (the weighted norm was unnecessary). Although the identification rates of SSEC were better than those provided by the RRC the method performed significantly worse than F-LR-IRNNLS. SSEC may not be effective in cases where faces are occluded by non-square objects.

The time performance for the block occlusion experiments is presented in Table I (columns 3 and 4)[5] for Yale B and AR datasets with the baboon image[6]. A key observation is that while F-IRNNLS achieved identical identification rates with RRC_L1 and RRC_L2 it is computationally more efficient by a magnitude.

B. Identification under Expressions & Face Disguise

In this experiment we tested our algorithms with face expressions and occlusion with real-world objects in three different scenarios: 1) faces with expressions 2) faces with sunglasses and 3) faces with scarves. In addition we wanted to evaluate our algorithms in cases where very few training instances are available per subject. Thus, the training set consists of the two neutral images (one from each session) from the AR dataset per subject. For the first scenario (face expressions) the 6 images per subject from sessions 1 and 2 with face expressions (smile, anger and scream) were selected for the testing set. For the second scenario (faces with sunglasses) the testing set consisted of the 6 images per subject with sunglasses from sessions 1 and 2. In the third scenario (faces with scarves) the 6 images per subject with scarves from sessions 1 and 2 were chosen for the testing set. The images were resized to 60×43 pixels. Identification rates for the three scenarios are shown in Table II for the various methods.

For the face expressions experiment all robust algorithms achieved high performance. A key observation for this experiment is that modeling the error as low-rank does not improve the results since face expressions errors do not form a contiguous area.

For the sunglasses experiment we observed that our F-LR-IRNNLS algorithm outperformed previous methods and was able to detect the outliers effectively. In this case modeling the error to be low-rank was adequate. This is due to the fact that the residual image consisted mainly of the sunglasses that made a contiguous error. SSEC performed poorly, perhaps, because the method does not capture well contiguous areas that are not square. A similar conclusion was drawn above with the random block occlusion experiment with a non-square image.

Results with the scarves experiment demonstrate that all methods robust to contiguous errors performed well, as expected (since the scarf occlusion is contiguous). Our F-LR-IRNNLS method achieved the best performance with 78.83% identification rate while our non-contiguous F-IRNNLS method achieved only 53.67%. This result emphasizes the fact that exploiting the spatial correlation in contiguous variations, such as scarves, is beneficial. Time performance is not reported here since all methods run very fast (less than a second) due to the fact that the training dictionary in this experiment was relatively small (200 training samples).

[5]In Table I we do not report results from the non-robust methods SRC and CR-RLS since they are both faster than all robust methods. However, they perform poorly in terms of identification rates in all experiments. In this work, our scope is to compare time performance between robust methods.

[6]Similar time results were obtained for the non-square image.

TABLE I: Average run time per test sample on the Extended Yale B and AR datasets under different variations.

Dataset	**Yale Baboon Occl. 60%**		**AR Baboon Occl. 50%**		**Yale Corruption 90%**		**AR Corruption 70%**	
	Accuracy	Time	Accuracy	Time	Accuracy	Time	Accuracy	Time
LR^3 [17]	88.57%	0.06s	49.86%	0.02s	14.29%	0.06s	n/a	n/a
HQ [11]	74.95%	3.11s	59.00%	0.99s	42.64%	1.96s	85.54%	1.49s
CESR [12]	41.76%	0.78s	54.86%	0.34s	36.04%	0.99s	85.69%	0.38s
RRC_L1 [10]	78.24%	12.41s	59.57%	3.95s	**84.40**%	15.23s	**91.92**%	29.84s
RRC_L2 [10]	81.54%	10.52s	55.57%	1.96s	76.26%	7.40s	88.38%	10.93s
SSEC [15]	76.70%	1.58s	59.29%	0.80s	4.84%	2.65s	n/a	n/a
Our F-IRNNLS	80.22%	1.52s	62.29%	0.45s	79.78%	2.84s	91.62%	2.25s
Our F-LR-IRNNLS	**95.82**%	2.41s	**74.43**%	0.57s	71.87%	4.30s	n/a	n/a

TABLE II: Identification Rates (%) under Face Disguise on the AR database.

Case	Expressions	Sunglasses	Scarves
SRC [4]	82.33%	37.17%	35.17%
CR-RLS [7]	81.33%	33.83%	39.33%
LR^3 [17]	79.50%	80.50%	77.00%
HQ [11]	94.33%	66.67%	45.67%
CESR [12]	93.50%	66.50%	17.17%
RRC_L1 [10]	94.33%	83.00%	58.50%
RRC_L2 [10]	93.67%	84.17%	72.50%
SSEC [15]	56.17%	70.67%	75.00%
Our F-IRNNLS	**94.83**%	81.50%	53.67%
Our F-LR-IRNNLS	93.00%	**89.83**%	**78.83**%

TABLE III: Identification Rates (%) and Time Performance (s) under Mixture Noise: Yale B 30% corruption & 60% occlusion, AR 20% corruption & 50% occlusion.

Dataset	**Yale B**		**AR**
	Accuracy	Time	Accuracy
SRC [4]	26.81%	1.04s	27.86%
CR-RLS [7]	14.73%	0.02s	28.29%
LR^3 [17]	44.62%	0.06s	35.43%
HQ [11]	42.20%	3.06s	42.00%
CESR [12]	23.96%	0.86s	40.86%
RRC_L1 [10]	43.08%	14.58s	47.00%
RRC_L2 [10]	41.54%	9.65s	38.14%
SSEC [15]	14.95%	1.72s	11.86%
Our F-IRNNLS	45.27%	1.45s	49.00%
Our F-LR-IRNNLS	**63.08**%	6.33s	**57.29**%

C. Identification under Mixture Noise

In this experiment we evaluate the performance of our algorithm for the case of mixture noise. In this case, both pixel corruption and block occlusion degraded the testing images. An example image with this degradation is shown in Figure 11. This experiment was conducted with two datasets, Extended Yale B and AR. Similarly to the previous Extended Yale B settings, Subsets 1 and 2 of Extended Yale B were used for training and Subset 3 was used for testing. With the AR dataset we chose the 700 non-occluded AR images for training from session 1 and the 700 non-occluded images for testing from session 2. In both datasets, for each testing image a percentage of randomly chosen pixels was corrupted. Corruption was performed by replacing those pixel values with independent and identically distributed samples from a uniform distribution between [0, 255]. Then, we placed the baboon square image on each corrupted test image. In Yale B dataset we performed this experiment with 30% pixel corruption and 60% occlusion. With the AR dataset, experiments were conducted with 20% pixel corruption and 70% occlusion. Identification rates are shown in Table III for the various methods.

F-LR-IRNNLS outperformed all previous methods which indicates that in the mixture noise case, our two error constraints capture the error term effectively. SSEC performed poorly due to the presence of pixel corruption. RRC_L1, RRC_L2 and HQ were robust to pixel corruption, however, their performance remained low since they were not effective on describing the occlusion part. Our F-LR-IRNNLS had a good balance on detecting the corrupted pixels and capturing the occlusion part with the employment of the weighted and nuclear norms. However, although F-LR-IRNNLS achieved significantly higher performance than the previous methods, the actual accuracy was relative low with 63.08% in YaleB and 57.29% in AR. The result may indicate that in mixture of noises further investigation about modeling the error is required.

Execution times in this case are reported in Table III. A key observation is that F-IRNNLS is by an order of magnitude faster than RRC_L1 and RRC_L2. F-LR-IRNNLS was faster than RRCs and slower than LR^3. However, LR^3 achieved significantly lower identification rates.

TABLE IV: Identification Rates on the Multi-PIE under Illumination.

Sessions	**Session 2**	**Session 3**	**Session 4**	
	Accuracy	Accuracy	Accuracy	Time
SRC [4]	95.48%	92.13%	95.71%	0.63s
CR-RLS [7]	95.30%	90.56%	94.46%	0.02s
LR^3 [17]	90.00%	83.44%	87.89%	0.04s
HQ [11]	95.84%	94.63%	97.14%	2.21s
CESR [12]	94.64%	92.94%	96.06%	1.94s
RRC_L1 [10]	95.84%	95.13%	97.14%	6.88s
RRC_L2 [10]	97.11%	94.19%	97.37%	31.11s
SSEC [15]	86.75%	76.13%	83.09%	5.40s
Our F-IRNNLS	**97.29**%	**96.00**%	**98.00**%	1.05s
Our F-LR-IRNNLS	96.87%	94.13%	97.83%	1.83s

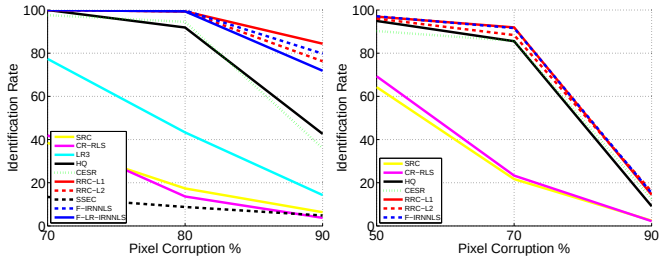

(a) Yale B dataset (b) AR dataset

Fig. 8: Identification Rates on the Extended Yale B and AR datasets under Pixel Corruptions.

D. Identification under Illumination

Experiments with variations in illumination were conducted on the Multi-Pie dataset. The Multi-PIE database [29] contains the images of 337 subjects captured in 4 sessions with simultaneous variations in pose, expression, and illumination. In the experiments we used all 249 subjects in Session 1. As in [7], we used 14 frontal images with 14 illuminations[7] and neutral expression from Session 1 for training, and 10 frontal images[8] from Sessions 2 to 4 for testing. Identification rates are shown in Table IV for the various methods.

Our first observation is that all methods achieved high identification rates. Simple SRC approaches performed well while robust methods only slightly improved the results. The reason is that in this experiment lighting intra-class variations exist in training samples and thus, the illumination variations are sufficiently reconstructed by the linear combination of the available training images.

With respect to time performance, our algorithm outperforms the previous robust methods. In particular the execution time in our approaches is around 1 second per test image while for RRC_L2 is around 30 seconds. Notice that although this was an experiment with a large training dictionary, our method retains very low running time.

E. Identification under Pixel Corruptions

Experiments under pixel corruption were conducted on two datasets: Extended Yale B and AR.

As in [4], [10] we used the non-occluded faces of Subsets 1 and 2 of the Extended Yale B for training (in total 719 images) and Subset 3 for testing (in total 455 images). Images were resized to 96×84 pixels. With the AR dataset, to make the experiment more challenging we chose to use occluded training and testing images. AR has 100 different subjects and for each subject the 13 images from session 1 (7 non-occluded images, 3 images with sunglasses, and 3 images with scarves) were used for training and the 13 images from session 2 for testing. Images were resized to 60×43 pixels.

For each test image in both datasets, a percentage of randomly chosen pixels was corrupted by replacing those pixel values with independent and identically distributed values from a uniform distribution between $[0, 255]$. The percentage of corrupted pixels was varied between 50 percent and 90 percent. Identification rates are shown in Figure 8 for the various methods.

Figure 8(a) illustrates that the robust non-contiguous methods RRC_L1 and F-IRNNLS achieved the best performance with over 80% accuracy in 90% pixel corruption. Methods able to handle contiguous errors such as SSEC, LR^3 and F-LR-IRNNLS performed poorly. We attribute this to the fact that pixel corruption is not a contiguous variation and modelling the error to have contiguous structure was inadequate. To that extend, with the AR dataset we decided to report results only on methods that handle non-contiguous errors as shown in Figure 8(b). In this dataset the accuracy is low in 90% corruption for all methods. As explained earlier, this was a more challenging experiment with a large number of testing images consisting of faces with pixel corruption on top of occlusion.

[7]Illuminations 0,1,3,4,6,7,8,11,13,14,16,17,18,19.
[8]Illuminations 0,2,4,6,8,10,12,14,16,18.

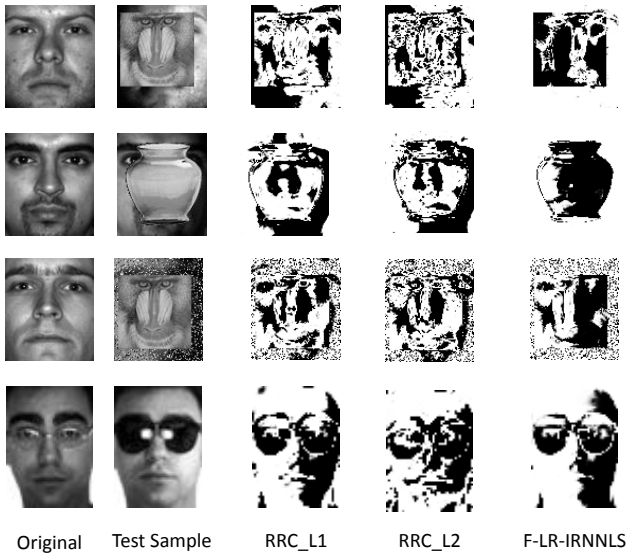

Fig. 9: Estimated weight maps for three iterative reweighted coding methods. The weight maps of our F-LR-IRNNLS method capture the outlier object of interest more accurately. With the other methods, a number of inlier pixels are detected as outliers.

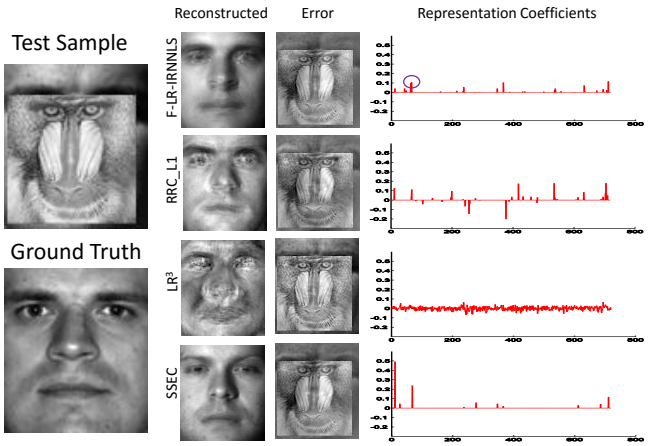

Fig. 10: Reconstruction results in 70% block occlusion.

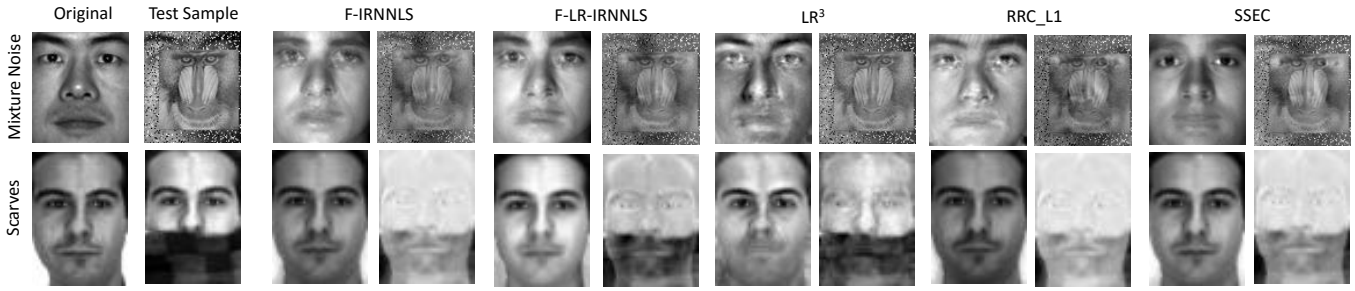

Fig. 11: Reconstruction results for various methods in mixture noise and scarves: The left image demonstrates the reconstructed face and the right image shows the estimated error.

With respect to execution time, our F-IRNNLS method outperformed RRC_L1 and RRC_L2 in both Yale B and AR datasets, as shown in Table I (columns 1 and 2). To emphasize the difference in performance, the execution time of F-IRNNLS in AR for 70% pixel corruption was about 2 seconds. In RRC_L1 and RRC_L2 was 29.84 and 10.93 seconds, respectively. Time performance of our algorithms was comparable to CESR and HQ (additive form) and worse than LR^3. However, these methods obtained lower identification rates.

Overall our proposed method achieved higher or competitive identification rates across all experiments. In addition, our method incurred lower computational costs than the previous algorithms. In some cases the execution time was lower by an order of magnitude than the second best algorithm overall (e.g., RRC_L1, RRC_L2). Furthermore, our F-LR-IRNNLS algorithm outperformed with respect to identification rates all previous methods with contiguous errors.

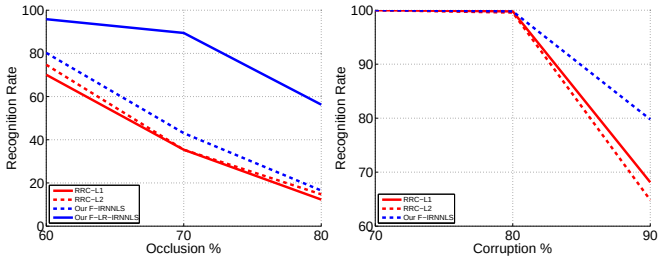

Dataset	Yale (60% occl.)		Yale (90% corr.)	
	Accuracy	Time	Accuracy	Time
RRC_L1 [10]	70.11%	9.16	68.13%	8.75
RRC_L2 [10]	74.73%	3.69	64.84%	3.77
Our F-IRNNLS	80.22%	**1.52**	**79.78%**	**2.84**
Our F-LR-IRNNLS	**95.82%**	2.41	n/a	n/a

Fig. 13: Performance of RRC ($t = 25$ iterations) and our method ($t = 100$ iterations) under Block Occlusion and Pixel Corruption.

F. Weight Map Estimations

Figure 9 shows the estimated weight maps between RRC_L1, RRC_L2 and our F-LR-IRNNLS in experiments with occlusions. Black values (close to zero) represent detected outliers by the various methods. We observe that F-LR-IRNNLS detected the outlier objects more effectively than the other methods. Most of the black regions in the weight maps are concentrated on the occluded area. In particular, we observe in Figure 9 (first row) that for the F-LR-IRNNLS method small weights are only assigned to the occluded (baboon) region as desired. On the other hand, the weight maps of RRC_L1 and RRC_L2 are not as accurate since outliers were detected in important pixels of the face. The reason is that with these methods there is no spatial correlation constraint between the weights. Similar conclusions can be drawn from all other examples in Figure 9.

G. Face Reconstruction Results

Figure 10 illustrates the face reconstruction results and the associated representation coefficients by the four methods. F-LR-IRNNLS and SSEC had the best reconstruction performance. The reconstructed face by LR^3 was poor mainly due to the choice of the regularizer for the representation coefficients (ℓ_2 norm). Similar reconstruction performance for the LR^3 method was encountered in almost all of our conducted experiments.

More reconstruction results for various methods with mixture noise and scarves variations are presented in Figure 11. With mixture noise, our F-LR-IRNNLS achieved the best reconstruction performance which demonstrates that our modeling was more effective in this case than the other methods. Face reconstruction was adequate for the case with scarves occlusion for all methods which validates the identification rates reported in Table II.

Figure 12 shows the (weighted) error and low-rank estimations during the ADMM iterations between F-LR-IRNNLS and LR^3 methods. From the error images we observe that our two-step approach in (22) (weighted and low-rank projections) estimates the error accurately during the ADMM iterations. In particular, the final error and reconstructed face are more accurately estimated by F-LR-IRNNLS than by LR^3.

H. Time Performance between our method and RRC

In this experiment we evaluate the identification rates in RRC [10] for the case where the maximum reweighted iterations $t = 25$. In other words, we want to investigate the performance degradation of RRC by keeping its execution time similar to our method. In our method we kept $t = 100$.

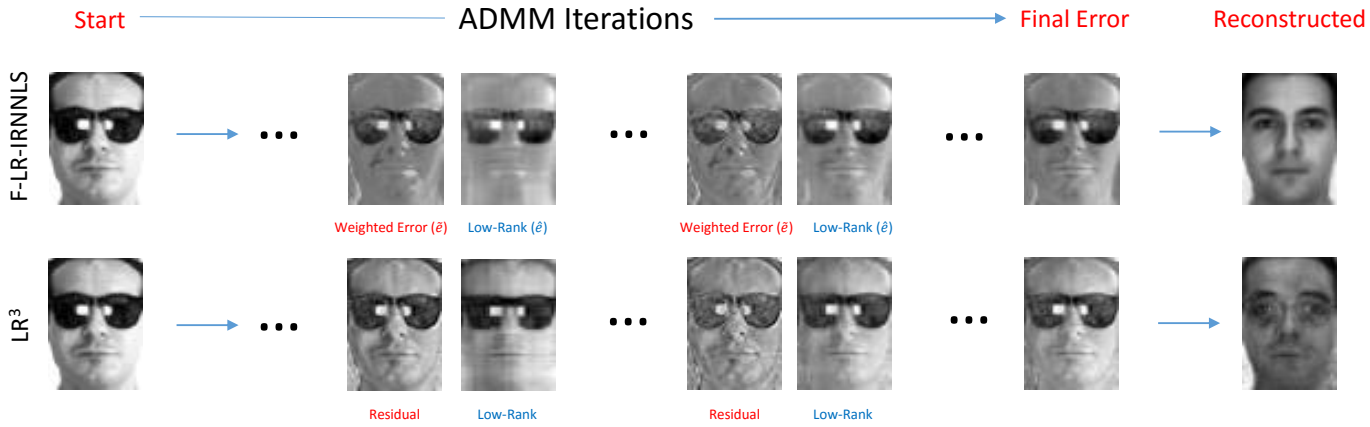

Fig. 12: Reconstruction results of the error and low-rank images estimated during the ADMM iterations.

TABLE V: Comparison of identification rates and times between RRCs and our Algorithms under different regularizations of the coefficients.

Dataset	**Illumination**		**Occlusion 60% (Y)**		**Occlusion 50% (A)**		**Corruption 70% (A)**		**Mixture Noise (Y)**	
	Accuracy	Time	Accuracy	Time	Accuracy	Time	Accuracy	Time	Accuracy	Time
RRC_L1 [10]	97.14%	6.88s	78.24%	12.41s	59.57%	3.95s	**91.92%**	29.84s	43.08%	14.58s
RRC_L2 [10]	97.37%	31.11s	81.54%	10.52s	55.57%	1.96s	88.38%	10.93s	41.54%	9.65s
F-IRLS	**98.06%**	1.32s	83.96%	1.74s	55.86%	0.39s	87.92%	1.17s	47.91%	1.70s
F-IRSC	97.77%	3.11s	79.78%	1.62s	60.43%	0.51s	90.08%	2.16s	45.05%	1.70s
F-IRNNLS	98.00%	1.05s	80.22%	1.52s	62.29%	0.45s	91.62%	2.25s	45.27%	1.45s
F-LR-IRLS	96.74%	1.36s	**99.12%**	2.81s	71.71%	0.56s	n/a	n/a	57.36%	9.38s
F-LR-IRSC	97.77%	2.79s	98.02%	2.50s	**74.57%**	0.68s	n/a	n/a	**66.81%**	7.96s
F-LR-IRNNLS	97.83%	1.83s	95.82%	2.41s	74.43%	0.57s	n/a	n/a	63.08%	6.33s

As shown in Figure 13, we observe that the computational time for RRC is now more competitive (although still higher than our method). However, the identification rates dropped significantly in both pixel corruption and block occlusion cases for RRC with $t = 25$.

I. Regularization of the Coefficients

In Table V we report performance comparisons of our method with different regularizations of the representation coefficients. Our main take away from the results is that sparsity is overall slightly better than the two other regularizers (non-negative and ℓ_2) in terms of identification rates. However, the non-negative regularizer provided a better balance between computational cost and identification rates. This is the reason it was preferred in our experiments.

Finally, one may observe the significant difference in time performance between the RRC and our method regardless of the regularization of the coefficients. The efficiency of our method gives rise to very large-scale robust face recognition systems for which computational time is a critical factor.

IV. Conclusions

In this work we proposed a method to describe contiguous errors effectively based on two characteristics. The first fits to the errors a distribution described by a tailored loss function. The second describes the error image as structural (low-rank). Our approach is computationally efficient due to the utilization of ADMM. The extensive experimental results support the claim that the proposed modeling of the error term can be beneficial and more robust than previous state-of-the-art methods to handle occlusions across a multitude of databases. A special case of our fast algorithm leads to the robust representation problem which is used to solve cases with non-contiguous errors. In this case, we showed that our fast iterative algorithm was in some cases faster by an order of magnitude than the existing approaches.

References

[1] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, "Face recognition: A literature survey," *ACM Comput. Surv.*, vol. 35, no. 4, pp. 399–458, Dec. 2003.

[2] M. Turk and A. Pentland, "Eigenfaces for recognition," *J. Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, Jan. 1991.

[3] P. Belhumeur, J. Hespanha, and D. Kriegman, "Eigenfaces vs. fisherfaces: recognition using class specific linear projection," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 19, no. 7, pp. 711–720, Jul 1997.

[4] J. Wright, A. Y. Yang, A. Ganesh, S. S. Sastry, and Y. Ma, "Robust face recognition via sparse representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 2, pp. 210–227, Feb. 2009.

[5] J. Wright, Y. Ma, J. Mairal, G. Sapiro, T. Huang, and S. Yan, "Sparse representation for computer vision and pattern recognition," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1031–1044, 2010.

[6] A. Wagner, J. Wright, A. Ganesh, Z. Zhou, H. Mobahi, and Y. Ma, "Toward a practical face recognition system: Robust alignment and illumination by sparse representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 2, pp. 372–386, Feb. 2012.

[7] D. Zhang, M. Yang, and X. Feng, "Sparse representation or collaborative representation: Which helps face recognition?" in *IEEE Int. Conf. Comp. Vision*, Nov. 2011, pp. 471–478.

[8] Q. Shi, A. Eriksson, A. Van den Hengel, and C. Shen, "Is face recognition really a compressive sensing problem?" *IEEE Conf. Comp. Vision Pattern Recognition*, vol. 0, pp. 553–560, 2011.

[9] M. Yang, D. Zhang, J. Yang, and D. Zhang, "Robust sparse coding for face recognition," in *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, June 2011, pp. 625–632.

[10] M. Yang, L. Zhang, J. Yang, and D. Zhang, "Regularized robust coding for face recognition," *Image Processing, IEEE Transactions on*, vol. 22, no. 5, pp. 1753–1766, May 2013.

[11] R. He, W.-S. Zheng, T. Tan, and Z. Sun, "Half-quadratic-based iterative minimization for robust sparse representation," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 36, no. 2, pp. 261–275, Feb 2014.

[12] R. He, W.-S. Zheng, and B.-G. Hu, "Maximum correntropy criterion for robust face recognition," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, no. 8, pp. 1561–1576, Aug 2011.

[13] R. He, W.-S. Zheng, B.-G. Hu, and X.-W. Kong, "Two-stage nonnegative sparse representation for large-scale face recognition," *Neural Networks and Learning Systems, IEEE Transactions on*, vol. 24, no. 1, pp. 35–46, Jan 2013.

[14] P. Huber, J. Wiley, and W. InterScience, *Robust statistics*. Wiley New York, 1981.

[15] X.-X. Li, D.-Q. Dai, X.-F. Zhang, and C.-X. Ren, "Structured sparse error coding for face recognition with occlusion," *Image Processing, IEEE Transactions on*, vol. 22, no. 5, pp. 1889–1900, May 2013.

[16] Z. Zhou, A. Wagner, H. Mobahi, J. Wright, and Y. Ma, "Face recognition with contiguous occlusion using markov random fields," in *Computer Vision, 2009 IEEE 12th International Conference on*, Sept 2009, pp. 1050–1057.

[17] J. Qian, J. Yang, F. Zhang, and Z. Lin, "Robust low-rank regularized regression for face recognition with occlusion," in *IEEE Conf. Comp. Vision Pattern Recognition Workshops*, Jun. 2014, pp. 21–26.

[18] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the Alternating Direction Method of Multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jan. 2011.

[19] S. Babacan, R. Molina, M. Do, and A. Katsaggelos, "Bayesian blind deconvolution with general sparse image priors," in *European Conference on Computer Vision (ECCV)*, September 2012, pp. 341–355.

[20] J. Palmer, K. Kreutz-Delgado, and S. Makeig, "Strong sub- and super-gaussianity," in *Latent Variable Analysis and Signal Separation*, ser. Lecture Notes in Computer Science, V. Vigneron, V. Zarzoso, E. Moreau, R. Gribonval, and E. Vincent, Eds., vol. 6365. Springer Berlin Heidelberg, 2010, pp. 303–310.

[21] R. Rockafellar, *Convex analysis*, P. U. Press, Ed., 1996.

[22] X. Zhou, R. Molina, F. Zhou, and A. Katsaggelos, "Fast iteratively reweighted least squares for lp regularized image deconvolution and reconstruction," in *Image Processing (ICIP), 2014 IEEE International Conference on*, Oct 2014, pp. 1783–1787.

[23] J. R. Shewchuk, "An introduction to the conjugate gradient method without the agonizing pain," Pittsburgh, PA, USA, Tech. Rep., 1994.

[24] A. Yang, A. Ganesh, S. Sastry, and Y. Ma, "Fast l1-minimization algorithms and an application in robust face recognition: A review," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2010-13, Feb 2010.

[25] S. Babacan, R. Molina, and A. Katsaggelos, "Bayesian compressive sensing using laplace priors," *Image Processing, IEEE Transactions on*, vol. 19, no. 1, pp. 53–63, Jan 2010.

[26] C. L. Lawson and R. J. Hanson, *Solving least squares problems*, ser. Classics in Applied Mathematics. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM), 1995, vol. 15, revised reprint of the 1974 original.

[27] A. Martínez and R. Benavente, "The AR face database," Computer Vision Center, Bellatera, Tech. Rep. 24, Jun. 1998.

[28] A. S. Georghiades, P. N. Belhumeur, and D. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 643–660, Jun. 2001.

[29] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multi-pie," *Image Vision Comput.*, vol. 28, no. 5, pp. 807–813, May 2010.

[30] D. L. Donoho and Y. Tsaig, "Fast solution of l1-norm minimization problems when the solution may be sparse." *IEEE Trans. Inf. Theory*, vol. 54, no. 11, pp. 4789–4812, 2008.